

CONTENT SWITCHING WITH USER-DEFINED POLICIES

Field of the Invention

[0001] This invention relates to computer based communication systems and more
5 particularly to content switching in communications systems.

Background

[0002] Basic computer based communications systems rely on a network
management entity to route and control the flow of traffic through the system. In
10 some cases nodes within the network have resources which determine a route from
the node to the next element in the connection. The routing of the communications
is based on the physical topology of the network and the policies determined by
the network operator. In packet-based networks, each packet contains a routing
information (the destination address in the case of IP networks), and the router
15 matches this routing information against its policy and routing topology
information to make the forwarding decision.

[0003] Content switches are currently being developed for web based
communications over the Internet or other IP-based communication networks.
20 These switches are known as being "content aware" and have the ability to
recognize the content being requested, where it is located, and how to route
individual content requests to the server best able to manage them. Basically, a
content switch is a router that makes routing or quality of service decisions based
on the content of the packets in the traffic flow rather than on pure routing
25 information. For example, a web switch, a content switch, is capable of
understanding the HTTP traffic and redirects it to the appropriate web server. A
classic application of such a system is a server farm, as shown in Figure 1, and will
be used as a reference throughout this application.

[0004] Web switches presently available are used to make switching decisions based on HTTP traffic. There is a trend, however, to use extensible mark up language (XML) over HTTP based protocols to provide web based applications such as Microsofts.Net framework. Simple object access protocol (SOAP) is an example of an XML over HTTP based protocol.

[0005] Active networks, in which the user has the ability to dynamically define routing policies, are also under development. A context type switch may be used in an active network wherein the full routing policy is provided along with the packets sent by the user. Generally, network providers are reluctant to embrace the active network framework in that the architecture lends itself to misuse by untrusted users.

[0006] The present invention addresses some of the problems associated with content switches. Current web switches have many drawbacks, one being that they are too dependent on the evolution of the protocol itself. For example, when the standard moved from HTTP1.0 to HTTP1.1 most content switches had to be replaced because they were examining only the first packets of the stream. Even in the context of a fully programmable web switch the evolution of the protocol has to be compensated by a programming effort which might be significant. As a more general note the content switch needs to support all of the applications switched on HTTP, H.323, RTP, etc. This is shown in Figure 2 which shows a typical protocol stack for a web switch. As mentioned earlier, web based applications are now based on XML making the classification at the HTTP layer insufficiently granular. For the sake of the present application the invention unifies the definition of the application under a common umbrella, for example XML.

[0007] Another problem with content switches is that the switching policy has to be defined by the application provider. This means, in practice, that web switches are

essentially controlled by the application provider. For active networks, the great flexibility that they allow is often regarded with some caution because of the trust issue that is implied by letting a user specify his own policies.

- 5 [0008] The present invention allows for a clear separation of the switching action and the switching policy.

Summary of the Invention

- 10 [0009] The present invention solves the problems of the prior art in as much as it provides a framework in which a content switch can be managed by a network provider while the applications that will be switched as well as the switching policies are defined by the application provider, a trusted customer of the network provider. The application provider will, typically, define switching policies responsive to individual customer requirements.

- 15 [0010] Therefore, in accordance with a first aspect of the present invention there is provided a schema stored on a computer-readable medium for use in a computer based communications system. The schema describes the structure of language in which the document is communicated; it includes as well routing rules, which, 20 when accessed to parse a document written in the language described in the schema, provide routing actions to be taken on packets belonging to a traffic flow associated with the document.

- 25 [0011] In accordance with a second aspect of the present invention there is provided a content switch for use in routing packets in a computer based communications system comprising: an entity capable of parsing a schema document associated with a packet and containing routing rules; and a routing instruction processor to interpret the routing rules in the schema document.

[0012] In accordance with a third aspect of the present invention there is provided a method of carrying out content switching in a computer-based communications system comprising: determining a routing action to be taken on packets of a flow associated with a document written according to a schema wherein the
5 determination is made by applying routing rules contained in the schema to elements of the document.

Brief Description of the Drawings

[0013] The invention will now be described in greater detail with reference to the
10 attached drawings wherein:

[0014] Figure 1 illustrates a content switching application involving a server farm; The administrative domains are circled: the web switch is in the same administrative domain than the servers running the application
15

[0015] Figure 2 shows an existing web switch protocol architecture; and

[0016] Figure 3 shows a content aware, content switch protocol architecture according to the present invention.
20

[0017] Figure 4 illustrates a similar application than Figure 1, this time using a XML-aware content switch and applications; the administrative domains of the XACS and the application servers are distinct.
25

Detailed Description of the Invention

[0018] The invention relates to a framework that consists of two main elements namely a content switch and a user-defined application that also contains a switching policy. It will be assumed that the present invention is applicable to one

architecture such as described in Figure 1 where the applications, running on the server farm, are based on XML. XML is a framework for defining mark up languages and is designed to separate syntax from semantics so as to provide a common framework for structuring information. In an XML-based application the structure of an XML document is described in its schema. The schema describes how to interpret the XML document.

[0019] A schema is a definition of the syntax of an XML based language i.e. a class of XML document and a schema language is a formal language for expressing schemas. Potentially, multiple schema languages are in use i.e. DTD, XSD, DSD etc. but for the purpose of the example, the W3C recommendation XML schema definition is used. The XML schema recommendation features processing instructions which can be used to pass specific information to the parser. These are in the form of XML tags such as `<?target data...?>`.

[0020] An XML parser sometimes referred to as the schema processor is used to check the validity of a document and, if the document is valid, extract the meaningful elements of the document. An XML document may contain reference to the schema it is using so that the schema does not have to be passed along with each document. The XML parser will get the document, retrieve the schema from the reference, and validate the document against the schema.

[0021] XML is used in this application as an example but it will be apparent to one skilled in the art that the concept is applicable to many other technologies (HTML CSS, SGML, etc), as described in section [0033].

[0022] The first element in the framework of the invention is a content switch which will also be known as an XML aware content switch. It is to be understood that the XML content switch is comparable to a usual web switch that is capable of

terminating HTTP sessions, and steering traffic to different servers. This will typically be done using network access translation (NAT) or other known techniques.

5 [0023] To satisfy the objectives of the invention, XML parsing capabilities are added to the content switch. To be able to parse an XML document, each document must contain a reference to its schema (a URI –Universal Resource Identifier- would be used). It will be assumed that the content switch can fetch the schema if necessary.

10

[0024] XML parsers capable of carrying out the parsing step are quite simple to realise because of the simplicity of the language itself. The XML parser will be connected to an XML routing instruction processor that is capable of interpreting the routing rules contained in the policy document. The protocol stack of an XML
15 aware content switch (XACS) is shown in Figure 3.

[0025] Having added parsing capabilities to the XACS the second piece of the invention is to add routing information to the schema. To avoid interfering with the description of the structure, routing information is added as a processing
20 instruction using specific tags, such as `<?xml_routing...?>`.

[0026] The tag contains a list of the routing/quality of service decisions that will be made by the content switch depending on the values of the associated element. These rules are ordered by priority and contain at least:

25

- an XML element to compare
- a value to compare with; and
- routing/QoS information if the element matches the value.

The rules might also contain a default action in the case the element in analysis does not match against the rules. In that case, if there is no default rule, the routing decision will be left to a usual routing engine (non content-aware).

5 [0027] An example of a schema with a routing rule is shown as follows:

```

10 <schema xmlns="http://www.w3.org/2001/XMLSchema"
    xmlns:b="http://www.mybusiness.com"
    targetNamespace="http://www.mybusiness.com">

    <element name="card" type="b:card_type"/>
    <element name="name" type="string"/>
    <element name="company" type="string"/>
    <element name="email" type="string"/>
15 <element name="phone" type="string"/>
    <element name="country" type="string"/>

    <complexType name="card_type">
        <sequence>
20     <element ref="b:name"/>
        <element ref="b:company"/>
        <element ref="b:email"/>
        <element ref="b:phone" minOccurs="0"/>
        <element ref="b:country" minOccurs="0"/>
25 </sequence>
    </complexType>

    <?xml_routing
    if (country = "France") {

```



```

    REDIRECT_LOAD_BALANCE (france1.mybusiness.com,
france2.mybusiness.com)
}
if (country = "Canada") {
5    REDIRECT (canada.mybusiness.com)
}
default:
    REDIRECT_LOAD_BALANCE (www1.mybusiness.com,
www2.mybusiness.com, www3.mybusiness.com)
10    ?>

</schema>

```

[0028] The language of the rules passed to the XML routing instruction processor
15 have to be fully defined but the definition thereof is not part of this invention.

[0029] As an example of a service using a content switch there is described an
application where a user has to fill in a form that contains his business card and
depending on the country where he lives, the contents which will redirect the
20 request to a different server. The aforementioned schema is associated to the
business card. The additions to the schema to support the invention are written in
bold print in the foregoing example.

[0030] The following shows a customer request for the business card:

```

<card xmlns=http://www.mybusiness.com
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.mybusiness.com/business_card.xsd">

```



```

<name>John Doe</name>
<company>Customer Business Inc.</title>
<email>john.doe@cbinc.com</email>
<phone>(202) 456-1414</phone>
5  <country>Canada</country>
    </card>

```

[0031] When a packet containing a customer request reaches the XML aware contents which the following happen

- 10 1. The XACS will send the XML document to the *XML parser*
2. The XML parser will fetch the schema from the schema location (http://www.mybusiness.com/business_card.xsd). A local cache can be added to speed up the process for subsequent access to the same schema.
3. The XML parser will validate the XML document according to the schema
- 15 4. The XML parser will pass the validated document to the `xml_routing` instruction processor to get the routing information
5. The `xml_routing` instruction processor will get the rule according to the schema (in this case the rule will be **REDIRECT canada.mybusiness.com**)
6. The XACS will redirect the packet to the server `canada.mybusiness.com` .

20

[0032] By letting the creator of the schema specify the routing policy, the application manager i.e. the creator of the schema, can define better services to its customers while relying on the network provider's contents switch to provide those services. This results in better quality of service and smarter routing, while
 25 separating the administrative domains of the content switch on one side and the application servers on the other side.

[0033] The concept works naturally with XML based applications because the technology itself contains the necessary enabler i.e. a schema.

While the description herein concentrates on XML it should be noted that:

- XML seems to be the standard of choice for future content-based application.
- Almost all the derivatives of SGML (as XML is) or any language that features a style sheet can be used in place of XML (which may make the parser more complex).
- Even HTML (starting with 4.0) features a style sheet that could be used for the purpose of the invention.
- A language that contains comments can also be used to provide such functionality, by putting the reference to the structure and the routing information in specially formatted comments.

[0034] So the invention is not limited to XML and is very applicable to most of the languages used in network-based applications.

[0035] By shifting the paradigm of content switching, the invention provides a framework in which a content switch can be managed by a network provider, while the applications that will be switched and the switching policies are defined by the application provider, a trusted customer of the network provider.

The trust relationship between the network and application providers can be described in a contract, and enforced by an entity in the XACS that checks whether the XML schema describing the routing decision respects the terms of the contract.

[0036] The invention enables high added value services to network and application providers and a better management of the cost it provides for:

- Better services (more added-value): the routing/QoS is done at the very fine application level.

- Reduced capital expenditures for the network provider: only one content switch for all the applications, and might be shared for multiple customers.
- Reduced operating expenditures for the network provider: the routing policy is managed by the application.
- 5 ◦ Reduced capital expenditures for the application provider: no content switch to buy
- Reduced operating expenditures for the application provider: the policy is now defined per application, each application provider directly managing the policy (instead of a central management on a single content switch).

10

[0037] Adding XML parsing capabilities creates an additional packet processing burden on the content switch. However, this represents a small portion of the overall processing that a content switch already performs.

15 [0038] Although particular embodiments of the invention have been described and illustrated, it will be apparent to one skilled in the art that numerous changes can be made without departing from the basic concept. It is to be understood, however, that such changes will fall within the full scope of the invention as defined by the appended claims.

20